A Bayesian Approach to High Resolution 3D Surface Reconstruction from Multiple Images

Bayesian Learning Group

NASA Ames Research Center

Peter Cheeseman, Esfandiar Bandari, Andre Jalobeanu, Frank Kuehnel, Robin Morris, John Stutz, Vadim Smelyanskiy, Doron Tal

Overview

- Computer Vision = Inverse Computer Graphics
- Bayesian Approach
  - General Solution to Inverse Problems
  - Hence *Bayesian Computer Vision*
- Theory
  - Light Scattering Model
  - Super-resolution
- Simplified Problem - 2D
- Initial 3D Problem
- Extended 3D Problem
- Summary

3D Surface Reconstruction - Aims

- **Build a high-resolution surface model that represents both the geometry and the reflectance properties of the surface using whatever image data is available.**
  - Useful for both science and navigation
  - integrate orbiter, lander/descent imagery and rover imagery
  - integrate new images into the existing model
  - integrate non-visual data (eg laser altimetery)

Computer vision = inverse graphics

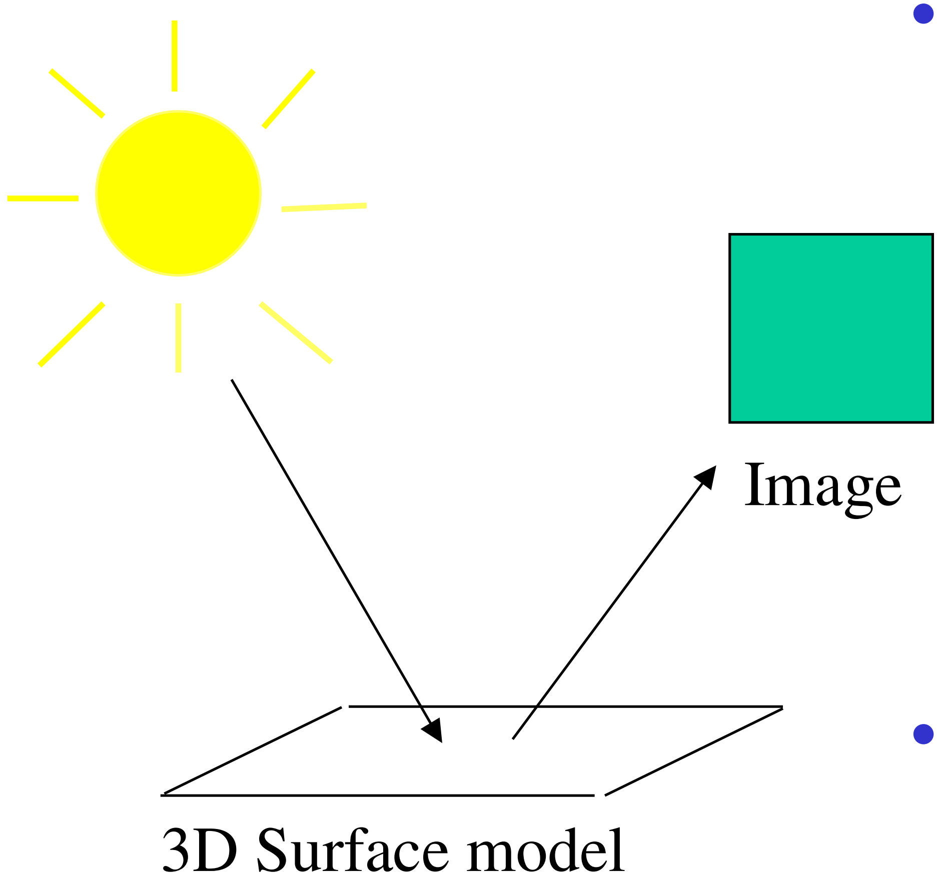

- Graphics
  - surface is known (to sufficient resolution)
  - lighting known
  - surface scattering properties known
  - camera parameters (position, fov, psf etc) known
- => compute *expected* images (pixels)

Computer vision = inverse graphics

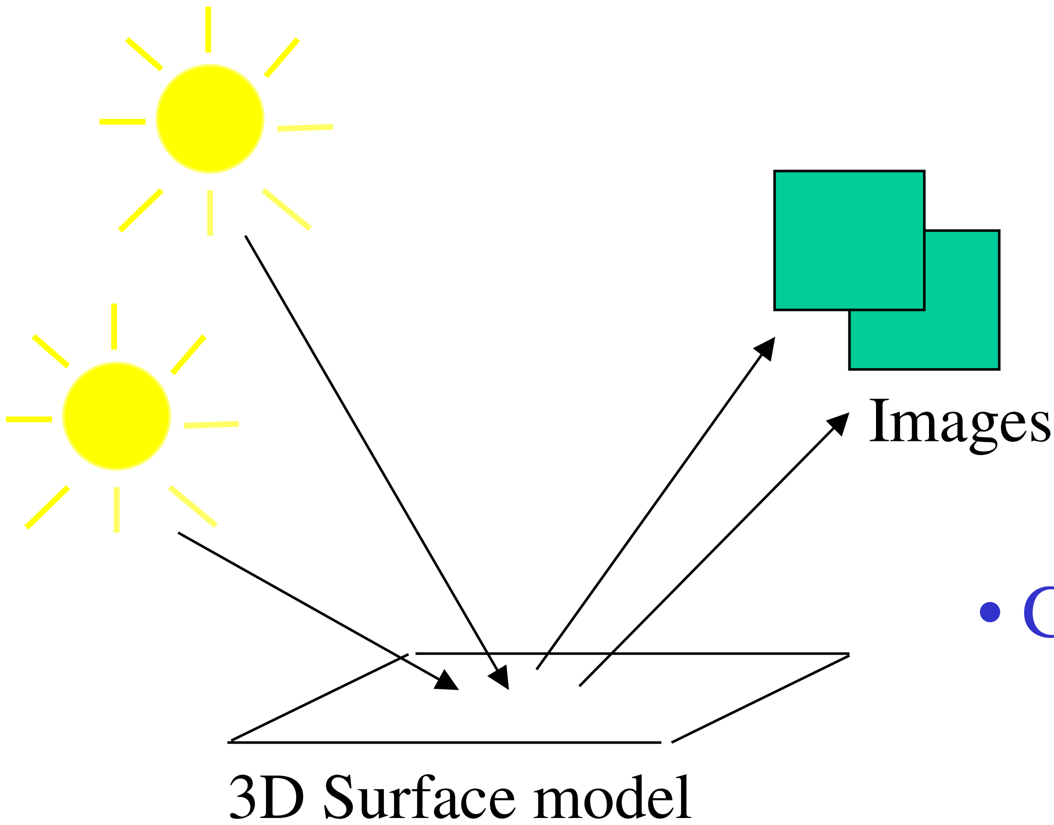

- Inverse Graphics
  - surface unknown
  - lighting unknown
  - surface scattering properties unknown (but scattering *model* known)
  - camera parameters unknown
- Given a set of images
  - find the most likely surface (and most likely values of other parameters)

Bayesian Solution to Inverse Problems

Ames Research Center

$$p(\text{3d surface} \mid \text{pixels}, \text{parameters}) \propto$$

$$p(\text{pixels} \mid \text{surface}, \text{parameters}) \times p(\text{surface} \mid \text{parameters})$$

$$\textit{likelihood} \qquad \times \qquad \textit{prior}$$

- parameters are
  - camera position and orientation, psf, field of view
  - lighting direction, strength
  - Initially parameters are assumed known. Will treat the unknown case later.

Likelihood

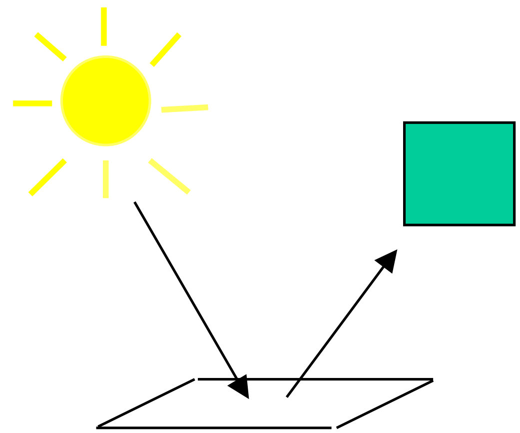

- Likelihood is

  $$p(\text{pixels} \mid \text{3D surface, parameters})$$

  – this is the graphics problem

  $$\text{Likelihood} = \prod_p \frac{1}{\sqrt{2\pi}\sigma_p} \text{Exp}[-\tfrac{1}{2}(\tfrac{I_p - \hat{I}_p}{\sigma_p^2})^2]$$

  – (independent gaussian model)

$\hat{I}_p$ = expected intensity of given pixel p

σ_p = standard deviation of actual pixel intensity I_p relative to $\hat{I}_p$

Putting the components together

- Prior
  - use a smoothness prior

$$p(h) \propto \exp(-h\Sigma^{-1}h^{T}/2)$$

- Posterior

$$-2\log P(\text{surface} \mid \text{pixels, parameters}) \propto$$

$$\tfrac{1}{\sigma_p^2}\sum_p (I_p - \hat{I}_p)^2 + h\Sigma^{-1}h$$

  - surface that maximizes the posterior = “regularized” least-squares estimate

3D Surface Model

- Triangulated surface
  - height field (hi)
  - regular grid (with local subdivision)
  - surface properties (eg albedo) associated with each triangle
- Advantages
  - guarantees surface continuity
  - compatible with existing graphics packages
- Alternatives
  - smoothly interpolated surface (splines)

Image Formation

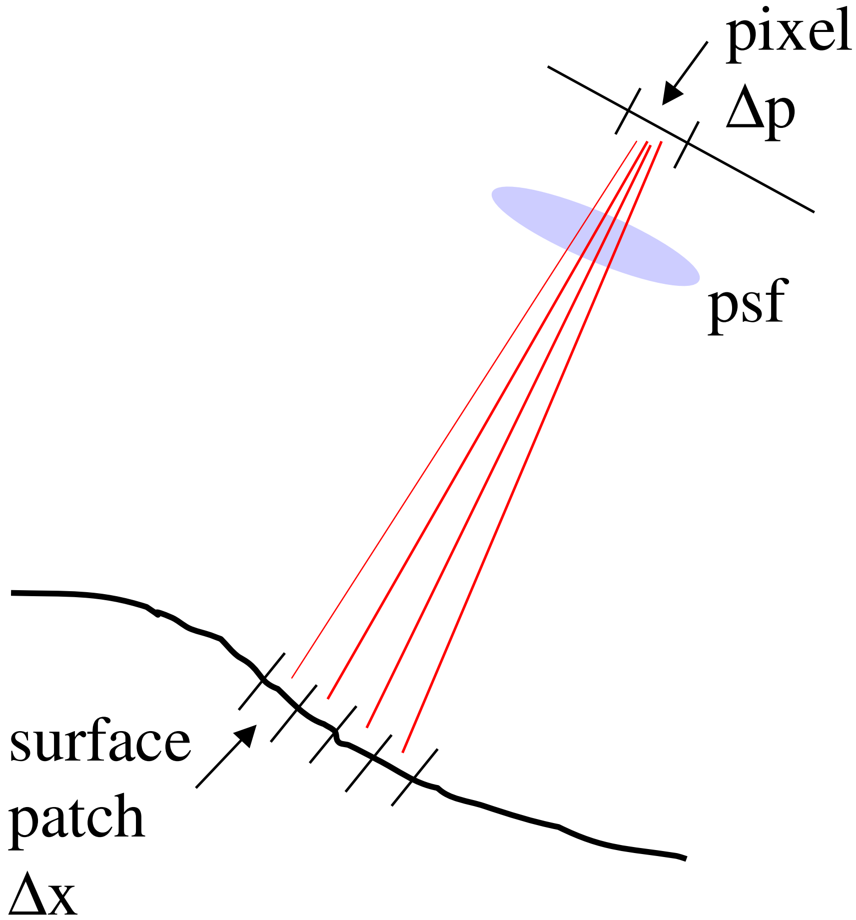

$\hat{I}_p$ – many to one mapping

– super - resolution,

$\Delta x < \Delta p$

$\hat{I}_p = \hat{I}_p$(lighting,

camera parameters

surface properties)

Comparison to standard rendering

Ames Research Center

- standard rendering
  - projected triangle >> pixel
- object-space rendering
  - projected triangle < pixel
  - needed for super-resolved inference

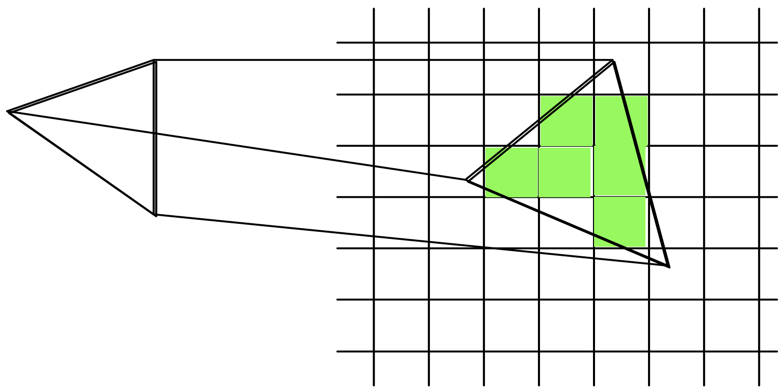

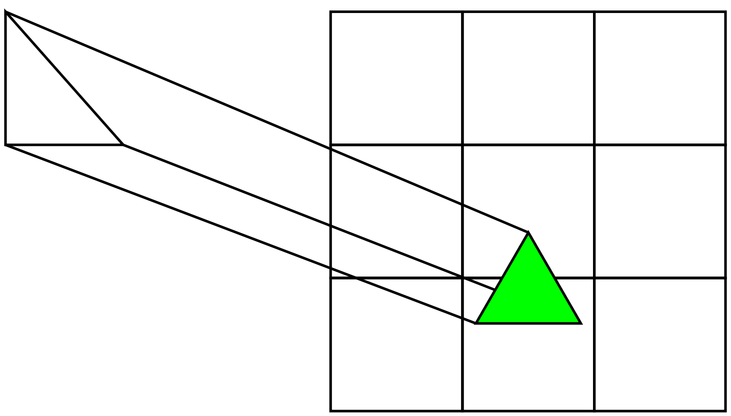

Why super-res works

- Beating the Nyquist Limit

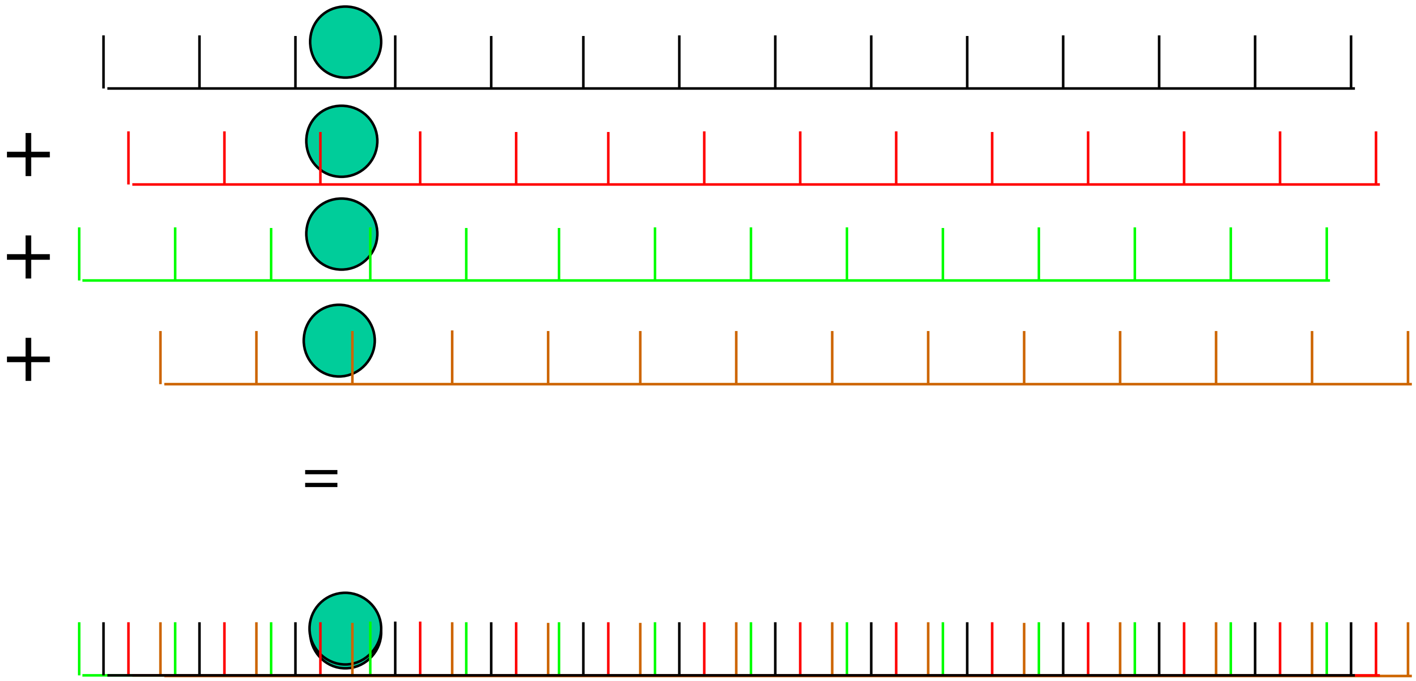

Surface Inference as Bayesian Inference

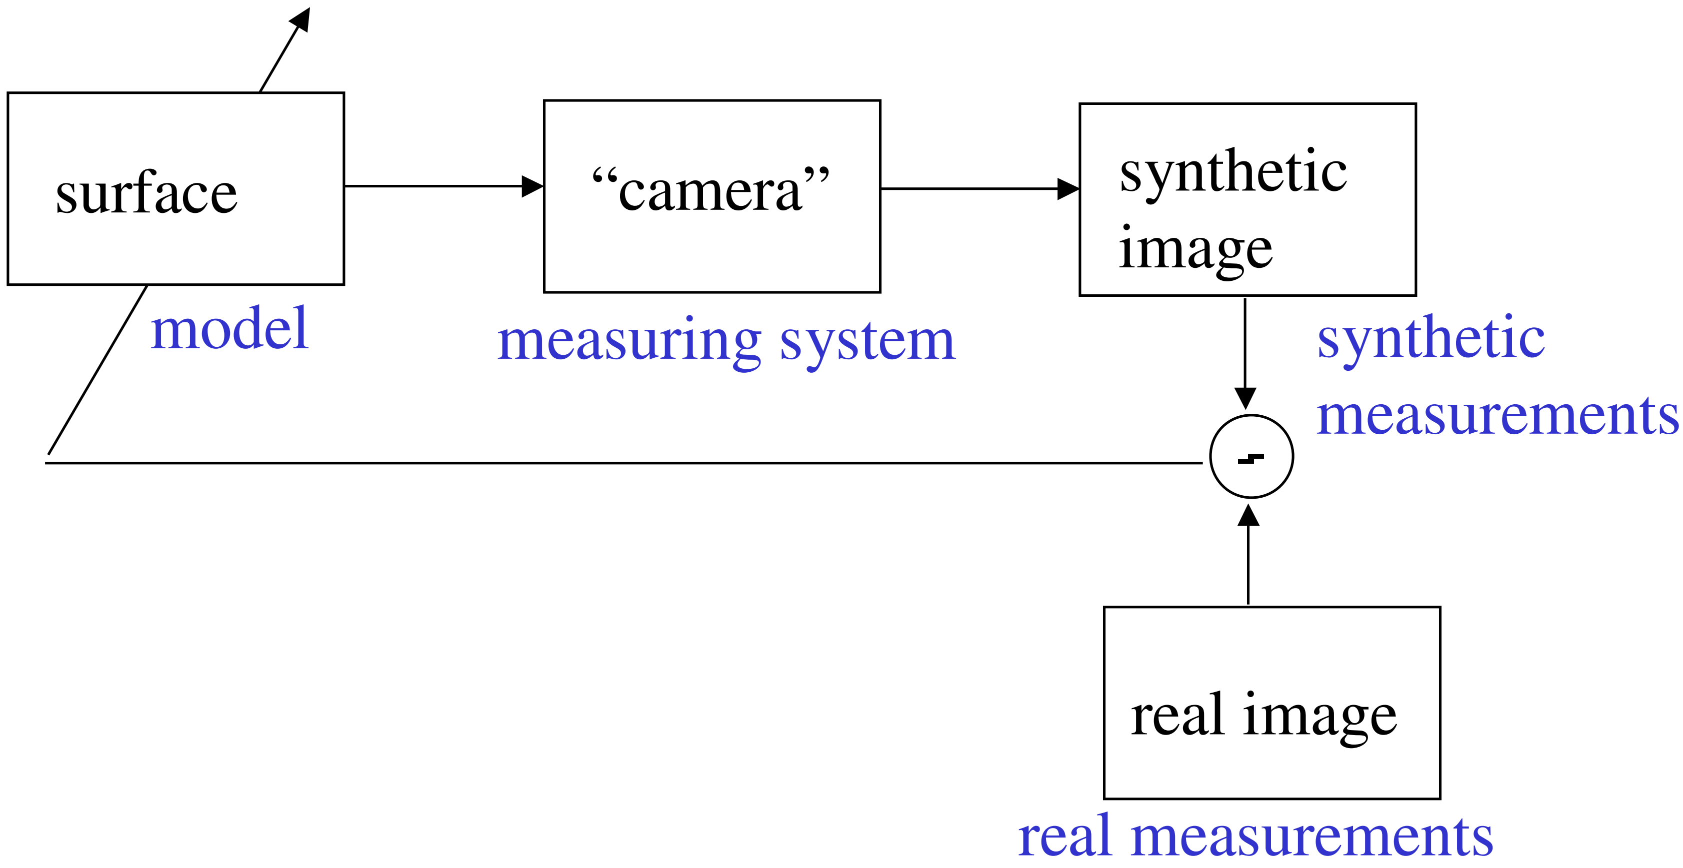

Bayesian Estimation of Surface Model Parameters

- Optimization
  - In the log-posterior, L(h,ρ) we linearize $\hat{I}(h,\rho)$ about the current estimate (h_0, ρ_0) giving

$$\hat{I}(h,\rho) \approx \hat{I}(h_0,\rho_0) + \mathbf{D}\begin{bmatrix} h-h_0 \\ \rho-\rho_0 \end{bmatrix}$$

  - Where

$$\mathbf{D}_{ij} = \frac{\partial \text{pixel}_i}{\partial \text{height(or albedo)}_j}$$

  - And the log-posterior is replaced by the quadratic form

$$L'(h,\rho) = \begin{bmatrix} h-h_0 \\ \rho-\rho_0 \end{bmatrix}^T \left(\Sigma^{-1} + \frac{\mathbf{DD}^T}{\sigma_e^2}\right)\begin{bmatrix} h-h_0 \\ \rho-\rho_0 \end{bmatrix} - \frac{(I-\hat{I}(h_0,\rho_0))}{\sigma_e^2}\mathbf{D}\begin{bmatrix} h-h_0 \\ \rho-\rho_0 \end{bmatrix} + \frac{\sum(I-\hat{I}(h_0,\rho_0))^2}{\sigma_e^2}$$

Error surfaces – heights and albedos

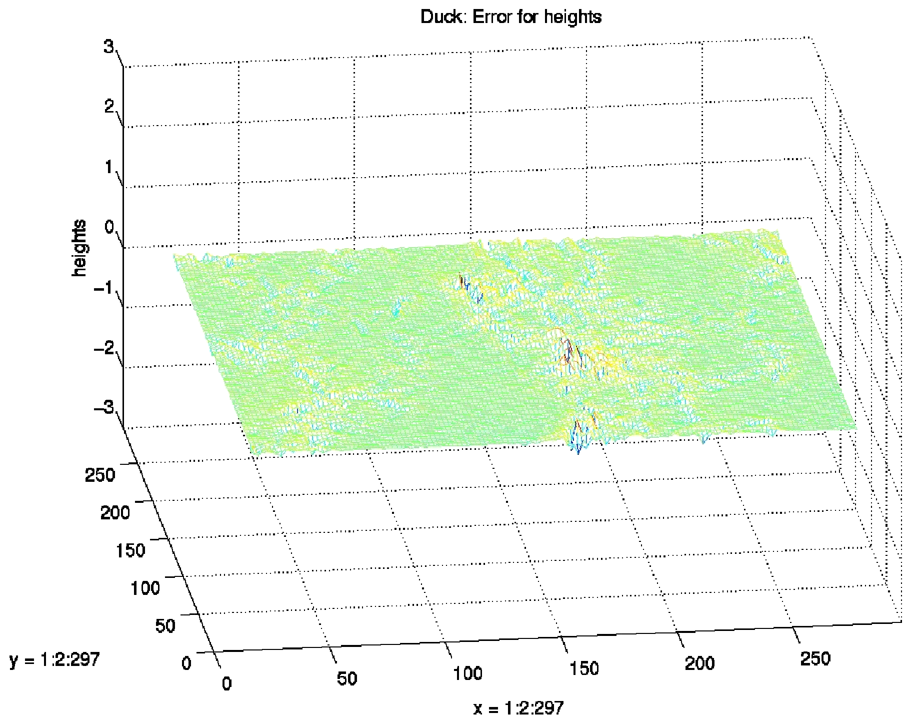

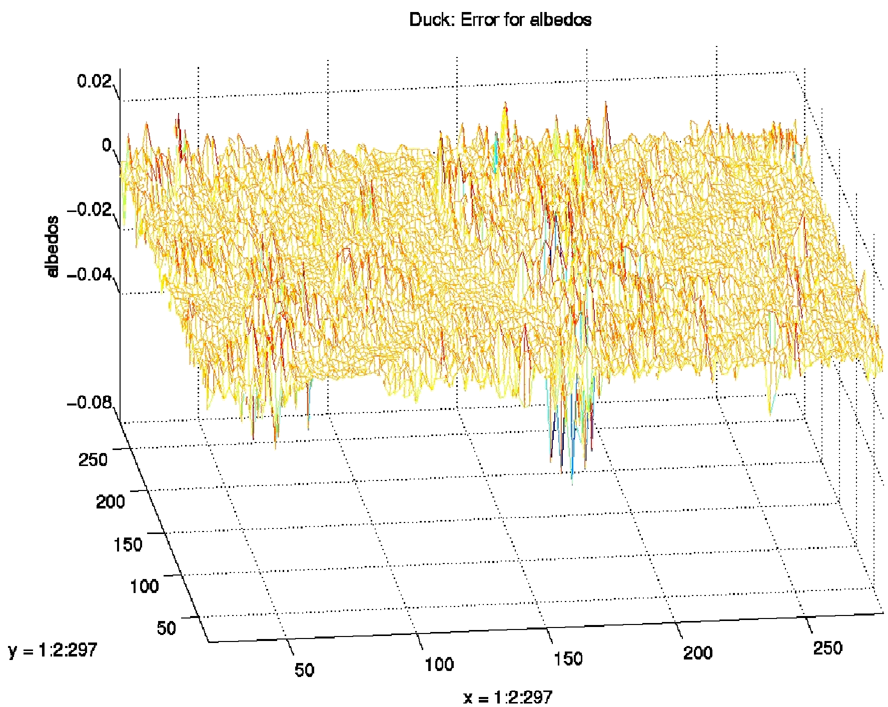

Inferred surface

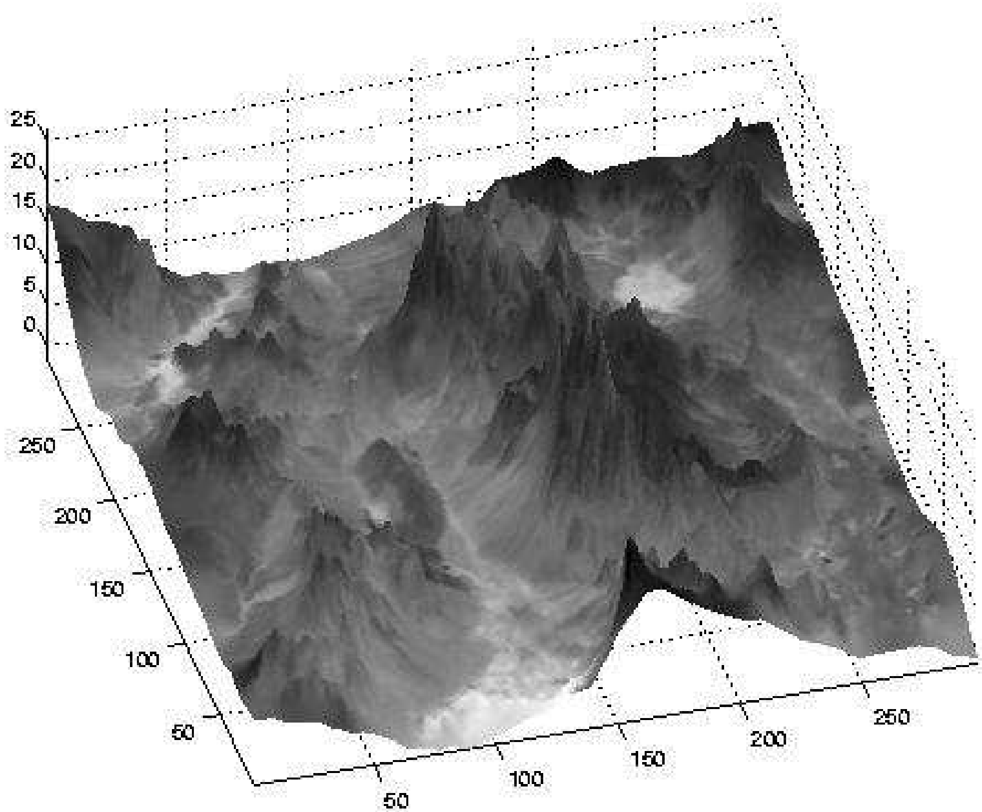

Real Input Images

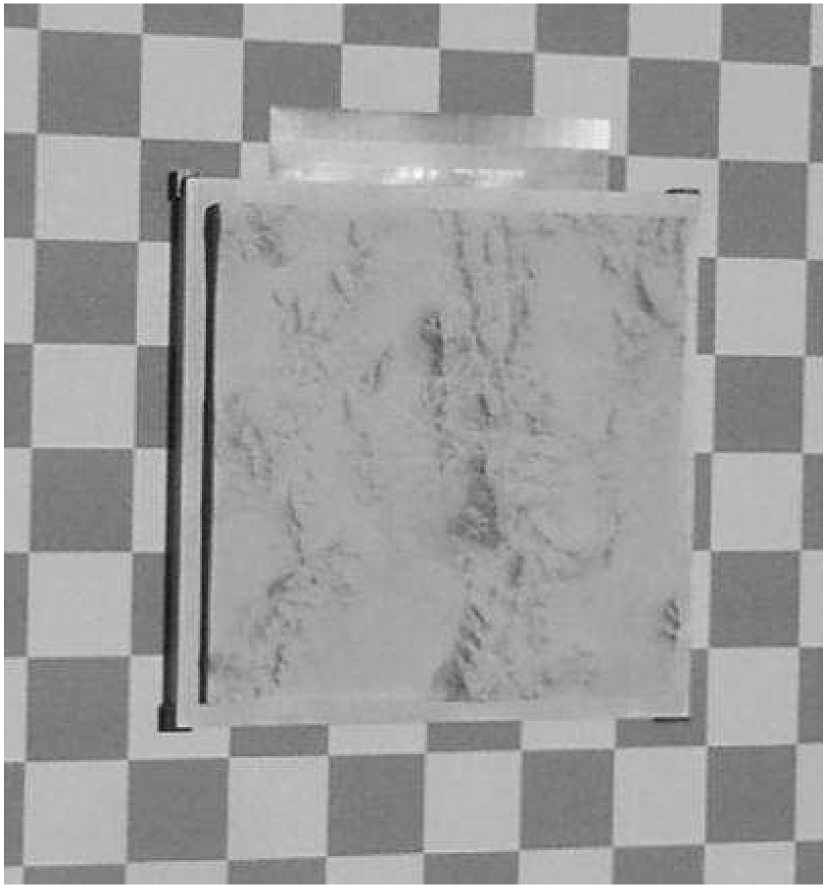

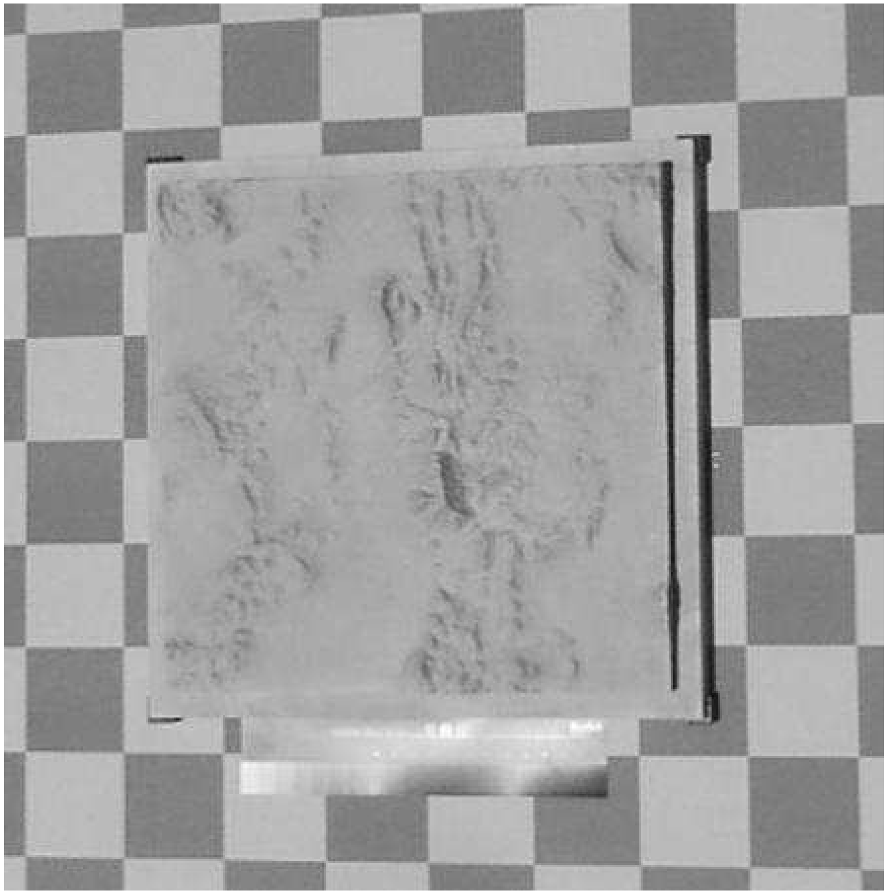

Reprojected Surfaces

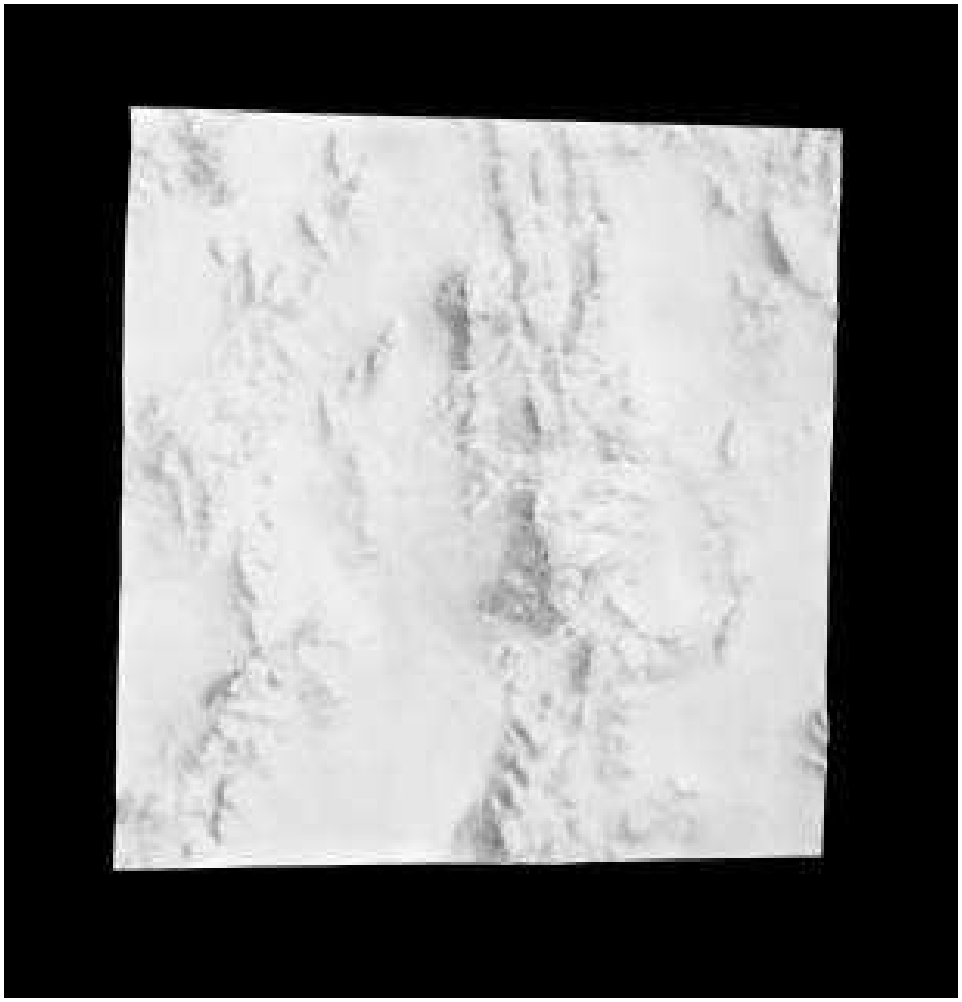

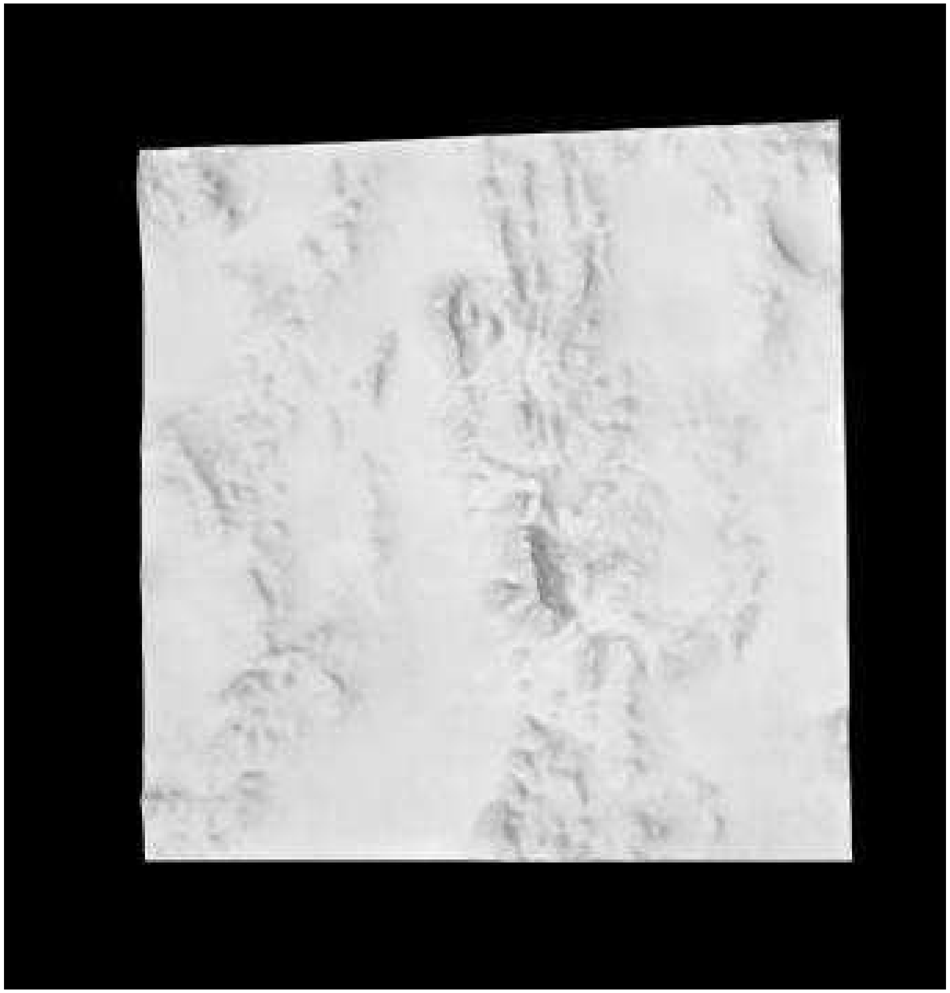

Summary

- Computer vision (3D surface reconstruction from multiple images) can be treated as *inverse graphics*
- Bayesian inference generally solves inverse problems, and can be applied to inverse graphics
- Reconstructed surface (and camera/lighting parameters) becomes a problem of “smoothed” parameter estimation (map estimation)
- Standard optimization procedures provide a practical solution to surface reconstruction
- Reconstruction can be at higher resolution than the images (super-resolution)